

CLAIMS

Sub
C1

1. A method of reading data from a remote memory of a remote device to a local memory of a local device across a network, said method comprising:

5 sending a message from said local device to said remote device, said message including a transport header indicating the message type of said message;

processing said message, in said remote device, to determine whether or not the transport header of said message identifies the message as a type of remote Direct Memory Access (rDMA) read operation; and

10 if the remote device determines from the transport header of said message that the message is said type of remote Direct Memory Access (rDMA) read operation, then performing a remote Direct Memory Access (rDMA) write operation to said local device in accordance with data elements included in said message.

2. The method recited in claim 1, wherein said data elements in said rDMA read message identify a set of source buffers in the remote device which reference the remote host-side memory and a set of destination buffers in the local device that reference the local memory.

Cont

5. The method recited in claim 4, wherein one data element of the rDMA read message specifies the last data segment and completion of the rDMA read request.

6. The method recited in claim 5, wherein the data is read from the remote memory of the remote device directly into the local memory of the local device over a virtual interface without making an intermediate copy of the data.

31

8. The method recited in claim 7, wherein the completion of the data transfer is processed in the local device, based on the immediate data that arrives with the last data segments of each rDMA write operation by the remote device.

9. A network device initiating a method to read data in the remote memory of a remote device directly into its memory, said network device having a network interface controller configured to carry out a method comprising:

sending a message from said local device to said remote device, said message including a transport header indicating the message type of said message;

processing said message, in said remote device, to determine whether or not the transport header of said message identifies the message as a type of remote Direct Memory Access (rDMA) read operation; and

if the remote device determines from the transport header of said message that the message is said type of remote Direct Memory Access (rDMA) read operation, then performing a remote Direct Memory Access (rDMA) write operation to said local device in accordance with data elements included in said message.

10. The network device recited in claim 9, wherein said data elements in said rDMA read message identify a set of source buffers in the remote device which reference the remote host-side memory and a set of destination buffers in the local device that reference the local memory.

11. The network device recited in claim 10, wherein the source buffers and destination buffers are registered with the Virtual Interface network interface controller of the remote device and the local device, respectively.

12. The network device recited in claim 11, wherein the data elements of the rDMA read message specify the source buffers and destination buffers as multiple data segments with offsets and designate a channel of the Virtual Interface as the data path for the rDMA write operation.

13. The network device recited in claim 12, wherein one data element of the rDMA read message specifies the last data segment and completion of the rDMA read request.

14. The network device recited in claim 13, wherein the data is read from the remote memory of the remote device directly into the local memory of

the local device over a virtual interface without making an intermediate copy of the data.

15. The network device recited in claim 14, wherein the remote device builds virtual interface rDMA write descriptors with the sequence inserted into the immediate data field on the last segment of each rDMA read request.

16. The network device recited in claim 15, wherein the completion of the data transfer is processed in the local device, based on the immediate data that arrives with the last data segments of each rDMA write operation by the remote device.

17. A tangible medium storing a plurality of program instructions, said program instructions causing a network device to carry out a method of reading data from the remote memory of a remote device directly to its local memory, said method comprising:

sending a message from said local device to said remote device, said message including a transport header indicating the message type of said message;

processing said message, in said remote device, to determine whether or not the transport header of said message identifies the message as a type of remote Direct Memory Access (rDMA) read operation; and

if the remote device determines from the transport header of said message that the message is said type of remote Direct Memory Access (rDMA) read operation, then performing a remote Direct Memory Access (rDMA) write operation to said local device in accordance with data elements included in said message.

15

C/Cont

18. The tangible medium recited in claim 17, wherein said data elements in said rDMA read message identify a set of source buffers in the remote device which reference the remote host-side memory and a set of destination buffers in the local device that reference the local memory.

19. The tangible medium recited in claim 18, wherein the source buffers and destination buffers are registered with the Virtual Interface network interface controller of the remote device and the local device, respectively.

20. The tangible medium recited in claim 19, wherein the data elements of the rDMA read message specify the source buffers and destination

buffers as multiple data segments with offsets and designate a channel of the Virtual Interface as the data path for the rDMA write operation.

21. The tangible medium recited in claim 20, wherein one data element of the rDMA read message specifies the last data segment and completion of the rDMA read request.

22. The tangible medium recited in claim 21, wherein the data is read from the remote memory of the remote device directly into the local memory of the local device over a virtual interface without making an intermediate copy of the data.

23. The tangible medium recited in claim 22, wherein the remote device builds virtual interface rDMA write descriptors with the sequence inserted into the immediate data field on the last segment of each rDMA read request.

Sub
A3 > 24. The tangible medium recited in claim 7, wherein the completion of the data transfer is processed in the local device, based on the immediate data that arrives with the last data segments of each rDMA write operation by the remote device.